

IN THE CLAIMS:

1 1. (Currently Amended) ~~In a file server having a storage operating system, a~~ method for
2 managing storage of data in a plurality of storage devices operatively connected to a
3 computer, each storage device having a plurality of blocks for storing data, comprising:

4 generating block layout information ~~in a file system layer of a~~ the storage
5 operating system executing on the computer by determining which blocks within the
6 plurality of blocks are allocated for storing data and which are unallocated;

7 ~~transferring the block layout information from the file system layer to a RAID~~
8 ~~layer of the storage operating system;~~

9 responsive to the block layout information, controlling execution of I/O
10 operations generated by the storage operating system at the RAID layer by identifying a
11 plurality of contiguous blocks ~~on a single storage device~~ within the plurality of blocks for
12 use by ~~the~~ each I/O operations so as to substantially maximize chain lengths of read
13 operations for calculation of parity;

14 determining whether a parity subtraction method or a recalculation method
15 requires a fewest number of read operations to calculate parity for the I/O operations;

16 selecting the parity subtraction method or the recalculation method for parity
17 calculation based on which method requires the fewest number of read operations; and

18 responsive to the block layout information and the parity calculation method
19 selected, identifying the contiguous blocks within the plurality of blocks for use by the
20 I/O operations.

1 2. (Currently Amended) A method for managing storage of data in a plurality of storage
2 devices operatively connected to a computer, each storage device comprising a plurality
3 of storage blocks, comprising:

4 generating block layout information of a storage operating system executing on
5 the computer by determining which blocks within the plurality of storage blocks are
6 allocated for storing data and which are unallocated;

determining whether a first methodology or a second methodology requires a
fewest number of read operations to calculate parity for I/O operations generated by the
storage operating system; and

in response to the block layout information and the determination, controlling
execution of I/O operations by identifying a plurality of contiguous storage blocks of the
plurality of storage blocks on a single storage device for use by ~~the~~each I/O operations so
as to ~~substantially~~ minimize ~~the~~ number of read operations needed for calculation of
error correction parameters across a stripe disposed among the plurality of storage
devices.

3. (Original) The method of claim 2 wherein the calculation of error correction
parameters comprises the calculation of parity.

4. (Currently Amended) The method of claim 2 wherein ~~the calculation of parity~~
~~comprises selecting a parity calculation operation from a group consisting of a~~
subtraction method ~~is~~as the first methodology and a parity re-calculation method ~~is~~as the
second methodology.

5. (Currently Amended) The method of claim 2 wherein the identification of the plurality
of contiguous storage blocks for use ~~by~~in the I/O operations ~~substantially~~ maximizes a
chain length by ~~substantially~~ maximizing a number of blocks of the plurality of storage
blocks having a contiguous physical layout on ~~the a first~~a single storage device of the
plurality of storage devices.

6. (Currently Amended) The method of claim 2, further comprising:
identifying the plurality of the contiguous storage blocks for use ~~by~~in the I/O
operations so as to ~~substantially~~ maximize ~~the~~ chain length by ~~substantially~~ maximizing
a number of blocks of the plurality of storage blocks having sequential volume block
numbers (VBNs) associated with the plurality of storage blocks.

1 7. (Currently Amended) The method of claim 2, further comprising:

2 identifying the plurality of the contiguous storage blocks for use ~~by~~ in the I/O
3 operations so as to ~~substantially~~ maximize ~~at~~ the chain length by ~~substantially~~ maximizing
4 locality of the plurality of contiguous blocks of ~~at~~ the first single storage device of the
5 plurality of storage devices.

1 8. (Currently Amended) The method of claim 2 wherein controlling execution comprises:

2 examining the plurality of storage blocks to which data is to be written ~~prior to~~
3 ~~write operations~~; and
4 selecting either,
5 the first methodology comprising minimizing a number of blocks read, or
6 the second methodology comprising maximizing chain lengths of blocks read for
7 calculating parity.

1 9. (Currently Amended) The method of claim 8, wherein controlling execution further
2 comprises:

3 implementing the selection of the first or second ~~parity calculation~~ methodology
4 responsive to the block layout information;
5 wherein, if the first methodology is selected ~~selection constitutes minimizing the~~
6 ~~number of blocks read~~,
7 determining on a stripe-by-stripe basis whether to calculate parity based on a
8 subtraction method or a recalculation method,
9 performing ~~a number of any appropriate~~ read operations of the plurality of storage
10 blocks to support the method selected, and
11 calculating parity responsive to the number of blocks read and the data to be
12 written; and
13 wherein, if the second methodology is selected ~~selection constitutes maximizing~~
14 ~~chain lengths of blocks read~~,
15 deciding which storage blocks to read to maximize chain length while minimizing
16 the number of blocks read to support either the subtraction method or the recalculation

17 method,
18 performing read operations on the number of blocks read, and
19 calculating parity responsive to the number of blocks read and the data to be
20 written.

1 10. (Currently Amended) The method of claim 2, wherein identifying the contiguous
2 storage blocks is based at least in part on an available resource of the computer.

1 11. (Currently Amended) The method of claim 2 further comprising transmitting the
2 block layout information from a file system layer of the computer to a RAID layer of the
3 computer.

1 12. (Currently Amended) The method of claim 2 wherein generating further comprises:
2 making a first determination as to whether a storage block of the plurality of
3 storage blocks is unallocated;
4 making a second determination as to a current implementation of the plurality of
5 storage devices; and
6 generating the block layout information based at least in part on the first and the
7 second determinations.

1 13. (Currently Amended) The method of claim 2, wherein ~~the I/O operation is one of a~~
2 ~~plurality of I/O operations and~~ at least one of the ~~plurality of I/O operations~~ is comprises
3 a read operation.

1 14. (Previously Presented) The method of claim 5, wherein the chain length is a length of
2 a read operation for calculation of parity.

1 15. (Previously Presented) The method of claim 5, wherein the chain length is a length
2 for a write operation for the data.

1 16. (Currently Amended) A method for managing storage of data in a storage system,
2 comprising:

3 maintaining a plurality of storage devices of the storage system, each storage
4 device having a plurality of storage blocks; and

5 writing data to predetermined storage blocks of the plurality of storage blocks
6 across a plurality of stripes and to predetermined contiguous storage blocks within the
7 plurality of each storage devices so as to maximize chain lengths of the predetermined
8 contiguous storage blocks ~~within each storage device~~ and minimizing a number of read
9 operations for the calculation of error correction parameters across ~~each stripe of the~~
10 plurality of stripes by determining whether a parity subtraction method or a recalculation
11 method requires a fewest number of read operations to calculate parity, and selecting the
12 parity subtraction method or the recalculation method for parity calculation based on
13 which method requires the fewest number of read operations.

1 17. - 38. (Cancelled)

1 39. (Currently Amended) A storage system, comprising:

2 a storage adapter configured to couple the storage system to a plurality of storage
3 devices, each storage device having a plurality of storage blocks; and

4 a storage manager in communication with the plurality of storage devices, the
5 storage manager configured to write data to predetermined storage blocks of the plurality
6 of storage blocks across a plurality of stripes and to predetermined storage blocks within
7 the plurality of each storage devices so as to ~~substantially~~ maximize a chain length of the
8 plurality of storage blocks by selecting ~~as many~~ contiguous storage blocks within a
9 first single storage device of the plurality of storage devices while minimizing a number
10 of read operations required for calculation of error correction parameters across ~~each~~
11 ~~stripe of the plurality of stripes~~ by determining whether a parity subtraction method or a
12 recalculation method requires a fewest number of read operations to calculate parity and
13 selecting the parity subtraction method or the recalculation method for parity calculation
14 based on which method requires the fewest number of read operations.

40. (Currently Amended) A system for managing ~~the data storage of data~~, comprising:
a plurality of storage devices operatively connected to a computer, each storage device having a plurality of storage blocks;
a storage device manager of the computer in communication with the plurality of storage blocks;
a block layout information generator of the computer in communication with the storage device manager and the plurality of storage blocks; and
an error correction parameter calculator of the computer in communication with the plurality of storage blocks and the storage device manager,
wherein the storage device manager, in response to ~~the~~ block layout information from the block layout information generator, controls execution of an I/O operation by identifying a plurality of contiguous storage blocks on one or more ~~a single~~ storage devices of the plurality of storage devices for use by the I/O operation so as to maximize a chain length within the one or more ~~single~~ storage devices while minimizing a number of read operations required ~~for calculation~~ by the error correction parameter calculator to calculate ~~of~~ error correction parameters across a stripe of the one or more storage devices by determining whether a parity subtraction method or a recalculation method requires a fewest number of read operations to calculate parity for the I/O operations and to selecting ~~the~~ the parity subtraction method or the recalculation method for parity calculation based on which method requires the fewest number of read operations.

41 - 44. (Cancelled)

45. (Currently Amended) A method for managing storage of data by a ~~server~~ computer, comprising:
receiving a request to write the data to a plurality of storage devices connected to the computer;
generating block layout information to determine which blocks within a plurality of blocks ~~located in~~ of the plurality of storage devices are allocated and which are unallocated;

identifying one or more blocks ~~of within~~ the plurality of blocks from the block layout information for use by a set of I/O operations;
determining a first number of read operations needed to calculate parity for the data by calculating parity using a subtraction method;
determining a second number of read operations needed to calculate parity for the data by calculating the parity using a recalculation method;
~~choosing either the subtraction method of calculating parity or the recalculation method of calculating parity by~~ determining which method requires a fewer number of read operations, and choosing the ~~which~~ method that requires the fewer number of read operations; and
writing the data to the identified one or more blocks, and calculating the parity for the data using the ~~which~~ chosen method.

46. (Currently Amended) The method of claim 45, further comprising:

choosing to either maximize chain lengths ~~of read operations~~ for calculating the parity or choosing to place the data with a high degree of locality in the plurality of storage devices for calculating the parity, ~~by choosing the which method that requires the fewer number of read operations.~~

47. (Currently Amended) A method for managing storage of data by a ~~server~~computer, comprising:

receiving a request to write data to a plurality of storage devices operatively connected to the computer;

generating block layout information to determine which blocks within a plurality of blocks ~~of located~~ in the plurality of storage devices are allocated and which are unallocated;

identifying the unallocated blocks ~~within the plurality of blocks~~ for use by a set of I/O operations to store the data; ~~and~~

determining, in response to the block layout information, whether a first method to minimize a number of read blocks or whether a second method to maximize chain

lengths of read blocks ~~based on which method~~ requires a fewer number of read operations, and

~~selecting implementing a selection one of the first method and the second method that requires the fewer number of read operations responsive to the determining and responsive to the block layout information, and during the writing of the data to the plurality of storage devices using the selected method, responsive to the block layout information, and responsive to whether minimizing the number of read blocks or substantially maximizing chain lengths of read blocks requires the fewer number of read operations.~~

48. (Currently Amended) The method of claim 47, further comprising:

in response to selecting the first method ~~to minimize the number of read blocks,~~ determining whether a subtraction method to calculate parity ~~based on a subtraction method~~ or whether a recalculation method to calculate the parity ~~based on which method~~ requires the fewer number of read operations; and

performing the write operation and calculating the parity using one of the subtraction method and the recalculation method ~~that which parity calculation method~~ requires the fewer number of read operations.

49. (Currently Amended) The method of claim 47, further comprising:

in response to selecting the second method ~~to maximize chain lengths of read blocks,~~ deciding which storage blocks to read to maximize chain lengths while minimizing the number of storage blocks read to support one of either the subtraction method to calculate parity and the a recalculation method to calculate the parity; and

performing the write operation and calculating the parity using one of the subtraction method and the recalculation method ~~which parity calculation method that~~ requires the fewer number of read operations.

50. (Currently Amended) A method for managing storage of data by a server computer, comprising:

receiving a request to write data to a plurality of storage devices operatively
connected to the computer;

generating block layout information to determine which blocks within a plurality
of blocks ~~located in~~ of the plurality of storage devices are allocated and which are
unallocated;

identifying one or more unallocated blocks ~~within the plurality of blocks~~ for use
by a set of I/O operations to store the data;

testing to either maximize chain lengths of read operations for calculation of
parity, or testing to place the data with a high degree of locality in the plurality of storage
devices, the testing comprising,

determining, for both maximizing chain lengths and placing the data with the high
degree of locality, a number of read operations needed to calculate parity for the data, by
calculating parity using both a subtraction method of calculating parity and a
recalculation method of calculating parity;

first choosing to either maximize chain lengths of read operations for calculation
of parity or to place the data with the high degree of locality in the plurality of storage
devices, and after the first choice, secondly choosing either the subtraction method of
calculating parity or the recalculation method of calculating parity by determining which
of these methods requires a fewest number of read operations,

choosing the method which ~~method~~ requires the fewest number of read operations
of calculating parity of the data; and

writing the data to the identified blocks, and calculating parity for the data using
the chosen ~~which~~ method.

51. (Currently Amended) A computer-readable media, comprising:

said computer readable media containing instructions for execution on a processor
for a method of managing storage of data in a plurality of storage devices, each storage
device having a plurality of blocks for storing data, the method comprising,

generating block layout information; and

in response to the block layout information, controlling execution of an I/O

7 operation ~~to the plurality of storage devices~~ by identifying a plurality of contiguous
8 storage blocks on ~~the plurality of a single~~ storage devices ~~for use by the I/O operation~~ so
9 as to minimize a number of read operations ~~needed~~ for calculation of error correction
10 parameters across a stripe of the plurality of storage devices by determining whether a
11 parity subtraction method or a recalculation method requires a fewest number of read
12 operations to calculate parity for the I/O operations and selecting one of the parity
13 subtraction method ~~and/or~~ the recalculation method for parity calculation based on which
14 method requires the fewest number of read operations.